

A multigrid solver for the 3D Monge-Ampère equation ^{*}

Jun Liu ^{‡||} Brittany D. Froese [‡] Adam M. Oberman ^{§||} Mingqing Xiao [¶]

August 26, 2015

Abstract

The elliptic Monge-Ampère equation is a fully nonlinear partial differential equation which has been the focus of increasing attention from the scientific computing community. Fast three dimensional solvers are needed, for example in medical image registration but are not yet available.

We build fast solvers for smooth solutions in three dimensions using a nonlinear full-approximation storage multigrid method. Starting from a second-order accurate centered finite difference approximation, we present a nonlinear Gauss-Seidel iterative method which has a mechanism for selecting the convex solution of the equation. The iterative method is used as an effective smoother, combined with the full-approximation storage multigrid method.

Numerical experiments are provided to validate the accuracy of the finite difference scheme and illustrate the computational efficiency of the multigrid algorithm. The solution time is almost linear in the number of variables. Problems of size 64^3 are solved in seconds and of size 128^3 are solved in a couple of minutes on a recent model laptop.

Keywords: Monge-Ampère equation; nonlinear Partial Differential Equations; finite difference method; Gauss-Seidel iteration; FAS multigrid method.

1 Introduction

The elliptic Monge-Ampère equation

$$\begin{cases} \det(D^2 u(x)) = f(x), & x \in \Omega \subset \mathbb{R}^d \\ u \text{ is convex} \end{cases}$$

^{*}Project was supported in part by NSF 1021203 of the United States.

^{||}The work described in this article is a result of a collaboration made possible by the IMA's workshop on "Topics on Control Theory: Optimal Mass Transportation" June 2014

[‡]Department of Mathematics and Statistical Sciences, Jackson State University, Jackson, MS 39217, USA. (jun.liu@jsums.edu)

[‡]Department of Mathematical Sciences, New Jersey Institute of Technology, University Heights, Newark, NJ 07102, USA. (bdfroese@njit.edu)

[§]Department of Mathematics and Statistics, McGill University, Montreal, QC H3A 0B9, Canada. (adam.oberman@mcgill.ca)

[¶]Department of Mathematics, Southern Illinois University, Carbondale, IL 62901, USA. (mxiao@siu.edu)

with $f > 0$, where D^2u is the Hessian of the function u , is a fully nonlinear partial differential equation (PDE) which has been the focus of increasing attention from the scientific computing community [Tad12, FGN13]. It is the prototypical fully nonlinear elliptic PDE, with connections to differential geometry [TW08]. As tools for solving the equation have become more effective, the PDE has appeared in more applications, including the reflector design problem [PTTBvR⁺14], the optimal transportation problem [BFO14], and parameter identification problems in seismic signals [EF14]. In this paper, we focus on building fast solvers for the PDE in the three dimensional case. Our interest in this case stems from the image registration problem [HZTA04, HRT10, HPM⁺09, HM07], where current imaging resolution requires more effective solvers than are currently available.

There are two regimes to consider when solving the problem. In the regime where solutions are smooth, many approaches can lead to numerical methods which converge in practice. These include finite differences [CDF11, LR05, SWR11], finite elements [BGNS11, BN12, FN09, Nei13], spectral methods [SAK13], and Fourier integral formulations [ZPF10].

In the second regime, where solutions are singular, standard methods either break down completely or become extremely slow [BFO10]. In order to compute singular solutions, the theory of viscosity solutions can be used to build convergent monotone finite difference methods [BS91]. Monotone schemes require using wide stencils, which have a directional resolution parameter, in addition to a spatial discretization parameter. As a result, they are generally less accurate than second order in space. On the other hand, they are stable in the maximum norm, and can be solved by a (slow) forward Euler method, with an explicit CFL condition [Obe08]. Refinements of the monotone method result in faster solvers [FO11b] and higher accuracy [FO12]. The fast solvers employed were an exact Newton's method (with an explicitly computed Jacobian) and a direct linear solver in two dimensions [FO12]. The initial guess for Newton's method used one step of an iterative solver from [BFO10], which also required a single solution of a Poisson equation. Computations were performed in three dimensions, but the direct linear solvers were too costly for larger sized problems.

A significant challenge in computing solutions of the Monge-Ampère equation is the need to enforce an additional constraint that the solution be convex; without this, the solution is not unique. In two-dimensions, for example, the Monge-Ampère equation will typically have two solutions: one convex and one concave. The situation becomes more complicated at the discrete level. For a two-dimensional problem with N discretization points, a finite difference scheme will typically admit 2^N different solutions [FGN13]. In three-dimensions, the problem becomes still more complicated. In two-dimensions, the correct solution can be selected by choosing the solution that has a positive Laplacian. This is equivalent to selecting the smallest root of a quadratic equation. However, in three-dimensions, positivity of the Laplacian and the Monge-Ampère operator is not enough to guarantee convexity. Using monotone schemes, convexity can be enforced by separately computing the positive and negative parts of various second directional derivatives.

In this paper, we focus on fast solution methods for the three-dimensional Monge-Ampère operator,

with a given source function which is positive and continuous in a convex domain $\Omega \subset \mathbb{R}^3$,

$$\det(D^2 u(x, y, z)) = f(x, y, z) > 0, \quad \text{for } (x, y, z) \text{ in } \Omega \quad (1.1)$$

with a given Dirichlet boundary condition

$$u = g, \quad \text{on } \partial\Omega \quad (1.2)$$

and the additional global constraint

$$u \text{ is convex,}$$

which is necessary for ellipticity and uniqueness. We record the explicit form of the operator in three dimensions as follows

$$\det(D^2 u) = u_{xx}u_{yy}u_{zz} + 2u_{xy}u_{yz}u_{xz} - u_{xx}(u_{yz})^2 - u_{yy}(u_{xz})^2 - u_{zz}(u_{xy})^2.$$

We note that the convexity constraint is formally equivalent to the condition that the Hessian of the solution is positive definite,

$$D^2 u > 0.$$

We assume also that that boundary data g is consistent with the restriction of a convex function to the boundary. These conditions on the data, along with strict convexity of the domain, are enough to ensure that solutions are classical (twice continuously differentiable) [CNS84]. For simplicity of computations, we work with square domains, which allow for the possibility of singular solutions, however, using known solutions avoids this problem.

Without additional difficulties, the method developed in this paper can also be applied to the two-dimensional problem

$$u_{xx}u_{yy} - (u_{xy})^2 = f(x, y) \quad \text{in } \Omega. \quad (1.3)$$

The objective of this work is to build fast solvers for the three-dimensional elliptic Monge-Ampere equation. We focus on the Dirichlet problem in the smooth solution regime, where the simple narrow stencil finite difference methods can be used. The goal is to build an effective, fast, scalable solver that can be used in applications. It is not trivial to build a multigrid solver in this setting—naive implementations of multigrid will fail. For example, the commonly suggested strategy [TOS01] of performing a few Newton iterations to approximately solve the corresponding nonlinear scalar equations at each grid point would not lead to a convergent FAS multigrid solver for our problem. This is because such a method has no mechanism for selecting the correct, convex solutions. We introduce such a selection mechanism, which leads to an effective multigrid method. However, when the equation becomes degenerate (corresponding to non-strictly convex solutions, which can arise when the right hand side $f(x) = 0$), the method presented here breaks down. This is likely related to the loss of uniform ellipticity of the equation, and the loss of strict convexity of the solution [BFO10]. A possible work around is to combine monotone (or filtered) schemes with a multi grid method. For now, we restrict our attention to

smooth solutions, where the method presented in this paper is effective.

One approach is to use Newton’s method with an iterative solver [OR00]. For example, the Newton-Multigrid method arises when a linear multigrid algorithm is used to approximately solve the linearized Jacobian system. A disadvantage of Newton’s method is that it requires a good initial guess, which should typically be convex [LR05]. We instead choose to treat the non-linearity directly using the linear multigrid framework and coarse-grid correction, which leads to the most common nonlinear multigrid method—the full approximation scheme (FAS) [Bra77, BL11, BHM00, TOS01, Saa03]. For the FAS multigrid to be effective as an iterative method, it requires an effective smoother (or relaxation) scheme such as a nonlinear Gauss-Seidel iteration, which must eliminate the high frequency components of the approximation errors at the current fine level. Once the FAS V-cycle iteration is established, the full multigrid (FMG) technique, based on nested iterations, can be exploited to obtain a good initial guess for the fine-grid problem.

In this work we develop a FMG-FAS multigrid algorithm for the 3D Monge-Ampère equation. The method is based on a nonlinear Gauss-Seidel iteration that includes a mechanism for selecting the convex solution. As a noticeable advantage over the above mentioned Newton-Multigrid method, the FMG-FAS multigrid algorithm usually has better global convergence properties. This is particularly the case for the Monge-Ampère equation since the Newton-Multigrid approach does not have any way to select the correct, convex solution. The FMG-FAS multigrid typically has a lower memory requirement as well since there is no need to compute and store the Jacobian matrices [TOS01].

There are very few published works devoted to designing fast nonlinear solvers, particularly for larger problem sizes in three dimensions. The multigrid method was used as a preconditioner for a Newton-Krylov iteration in [CDF11, DCF⁺08] in two and three dimensions up to grids of size 128^3 , which took approximately five minutes using eight processors. For the applications considered in that work, solutions of the Monge-Ampère equation would typically be close to $u = |x|^2/2$, which led naturally to a good initialization of Newton’s method. In this paper, we approach the problem using the nonlinear full approximation scheme (FAS) multigrid method, which will not require an accurate initialization. Although the FAS multigrid methods are popular and widely used in many disciplines involving nonlinear PDEs, such as variational image registration models [CC11], the only references we found on applying FAS multigrid method to the Monge-Ampère equation were two-dimensional simulations in a master’s thesis [Soi11]. It is worthwhile to point out that the reported computational CPU time does not scale linearly with respect to the degrees of freedom (see e.g. Table 5.1 in [Soi11]), which motivates us to look further into the efficiency of the FAS multigrid implementations. In additions, we are mainly interested in solving three dimensional (3D) Monge-Ampère equation, which is well-known to be much more difficult [BN12].

This paper is organized as follows. In the next section, we present a discretization of the PDE using second-order centered finite differences and propose a nonlinear Gauss-Seidel iterative method for solving the discretized system. The method extends the two-dimensional Gauss-Seidel method of [BFO10]; in our case, we need to solve a cubic rather than a quadratic equation. In Section 3, we introduce a full multigrid method based on the FAS V-cycle algorithm, where the nonlinear Gauss-Seidel iterations func-

tion as a smoother. In Section 4, numerical results for several two- and three-dimensional examples are reported, which demonstrate the accuracy, mesh independent convergence, and linear time complexity of our proposed FMG-FAS multigrid solver. Finally, the paper ends with several concluding remarks in Section 5.

2 A nonlinear Gauss-Seidel iteration in 3D

This section is motivated by the two-dimensional explicit finite difference method in [BFO10] where second derivatives are discretized using standard centered differences on a uniform Cartesian grid. A Gauss-Seidel iteration is constructed by selecting the smaller root of a point-wise defined quadratic equation over each grid point. The resulting method is second order accurate provided the solution is regular enough. In the following, we will develop a similar Gauss-Seidel iteration for the three-dimensional case, where we have to handle a cubic equation at each grid point.

Let $N > 0$ be a positive integer. We discretize the space domain $\bar{\Omega} = [0, 1]^3$ uniformly into

$$\Omega_h = \{(x_i, y_j, z_k) \mid x_i = ih, y_j = jh, z_k = kh; i, j, k = 1, \dots, N\}$$

with a uniform mesh step size $h = 1/N$. Let $u_{i,j,k}$ be the discrete approximation of $u(x_i, y_j, z_k)$. Notice that the values of $u_{i,j,k}$ with $i \in \{0, N\}$ or $j \in \{0, N\}$ or $k \in \{0, N\}$ are directly specified by the Dirichlet boundary conditions. Hence, we only need to set up finite difference approximations for all the interior grid points. We employ the second-order accurate centered finite difference approximations [LeV07] for all the second-order derivatives as follows.

$$u_{xx}^h(x_i, y_j, z_k) = \frac{1}{h^2} (u_{i+1,j,k} + u_{i-1,j,k} - 2u_{i,j,k}) + O(h^2) \quad (2.1)$$

$$u_{yy}^h(x_i, y_j, z_k) = \frac{1}{h^2} (u_{i,j+1,k} + u_{i,j-1,k} - 2u_{i,j,k}) + O(h^2) \quad (2.2)$$

$$u_{zz}^h(x_i, y_j, z_k) = \frac{1}{h^2} (u_{i,j,k+1} + u_{i,j,k-1} - 2u_{i,j,k}) + O(h^2) \quad (2.3)$$

$$u_{xy}^h(x_i, y_j, z_k) = \frac{1}{4h^2} (u_{i+1,j+1,k} + u_{i-1,j-1,k} - u_{i-1,j+1,k} - u_{i+1,j-1,k}) + O(h^2) \quad (2.4)$$

$$u_{xz}^h(x_i, y_j, z_k) = \frac{1}{4h^2} (u_{i+1,j,k+1} + u_{i-1,j,k-1} - u_{i-1,j,k+1} - u_{i+1,j,k-1}) + O(h^2) \quad (2.5)$$

$$u_{yz}^h(x_i, y_j, z_k) = \frac{1}{4h^2} (u_{i,j+1,k+1} + u_{i,j-1,k-1} - u_{i,j+1,k-1} - u_{i,j-1,k+1}) + O(h^2). \quad (2.6)$$

Using these approximations, we can construct the discrete Hessian

$$D^2 u^h = \begin{pmatrix} u_{xx}^h & u_{xy}^h & u_{xz}^h \\ u_{xy}^h & u_{yy}^h & u_{yz}^h \\ u_{xz}^h & u_{yz}^h & u_{zz}^h \end{pmatrix}.$$

Then the discrete form of the elliptic Monge-Ampère equation in the interior of the domain reads

$$\begin{cases} \det(D^2 u_{i,j,k}^h) = f_{i,j,k} \\ D^2 u_{i,j,k}^h > 0. \end{cases} \quad (2.7)$$

Inserting the approximations into equation (2.7) at each interior grid point (x_i, y_j, z_k) leads to a cubic polynomial $P_3(u_{i,j,k}, u)$ with respect to $u_{i,j,k}$

$$\begin{aligned} P_3(u_{i,j,k}; u) := & -8u_{i,j,k}^3 + 4(a+b+c)u_{i,j,k}^2 + 2(r^2 + s^2 + t^2 - ab - ac - bc)u_{i,j,k} \\ & + (abc - at^2 - bs^2 - cr^2 + 2rst - h^6 f_{i,j,k}) = 0, \end{aligned} \quad (2.8)$$

where we use the notation

$$\begin{aligned} a &= (u_{i+1,j,k} + u_{i-1,j,k}), \\ b &= (u_{i,j+1,k} + u_{i,j-1,k}), \\ c &= (u_{i,j,k+1} + u_{i,j,k-1}), \\ r &= \frac{1}{4} (u_{i+1,j+1,k} + u_{i-1,j-1,k} - u_{i-1,j+1,k} - u_{i+1,j-1,k}), \\ s &= \frac{1}{4} (u_{i+1,j,k+1} + u_{i-1,j,k-1} - u_{i-1,j,k+1} - u_{i+1,j,k-1}), \\ t &= \frac{1}{4} (u_{i,j+1,k+1} + u_{i,j-1,k-1} - u_{i,j+1,k-1} - u_{i,j-1,k+1}), \end{aligned}$$

and $f_{i,j,k} = f(x_i, y_j, z_k)$.

The above finite difference approximations (2.8) are intentionally formulated as a cubic equation of $u_{i,j,k}$ at the current node (x_i, y_j, z_k) by assuming its neighboring nodes are fixed. The nonlinear Gauss-Seidel iteration follows from iteratively solving the cubic equation sequentially over all grid points in a lexicographic order. The so-called red-black ordering with better smoothing and parallel properties is usually recommended for 2D problems, but it is of limited advantage [Zha98] to our 3D scheme involving a total of 19 points. In this case, four or more colors are necessary to fully parallelize the updating order [GZ00], which is not further pursued in the current paper. The Dirichlet boundary conditions (1.2) are enforced at the boundary grid points during the iteration. However, we also need to determine which root of the cubic is to be selected in order to ensure local convexity of the discrete solutions. In the 2D case, the corresponding approach leads to a quadratic equation [BFO10], and selection of the smallest root yields the convex solution. This rule is not directly applicable for the cubic equation arising in 3D, which may have either one or three real roots. As we are only interested in real-valued solutions, we simply select the smallest of all the real roots; that is

$$u_{i,j,k}^{n+1} = \min \left\{ \text{real roots of } P_3(\cdot; u^n) \right\}. \quad (2.9)$$

This selection of the correct root is justified in Theorem 2.1. These roots can either be computed exactly or approximately using a root-finding algorithm.

Theorem 2.1. *Let u be a solution of the discrete Monge-Ampère equation (2.7). Then u is a fixed point of (2.9).*

Proof. Consider a fixed location (x_i, y_j, z_k) in the grid. Since $u_{i,j,k}$ satisfies the discrete Monge-Ampère equation, it automatically satisfies the cubic equation $P_3(u_{i,j,k}; u) = 0$ as well. It remains to show that $u_{i,j,k}$ is the smallest real root of this cubic.

One possibility is that the cubic equation has only a single real root, in which case this must coincide with the real-valued $u_{i,j,k}$.

The other option is that the cubic equation has three real roots, $v_1 \leq v_2 \leq v_3$. We remark that using the notation defined above, the discrete Hessian corresponding to each of these roots can be written as

$$D^2(u^h; v_m) = \begin{pmatrix} \frac{a^h - 2v_m}{h^2} & u_{xy}^h & u_{xz}^h \\ u_{xy}^h & \frac{b^h - 2v_m}{h^2} & u_{yz}^h \\ u_{xz}^h & u_{yz}^h & \frac{c^h - 2v_m}{h^2} \end{pmatrix}, \quad m = 1, 2, 3 \quad (2.10)$$

where the off-diagonal elements do not depend on the value of v_m . This is a symmetric real-valued matrix, and therefore has real eigenvalues.

Now suppose that λ is any eigenvalue of $D^2(u^h; v_1)$. Then we can compute

$$D^2(u^h; v_m)x = \left(D^2(u^h; v_1) + \frac{2(v_1 - v_m)}{h^2} I \right) x = \left(\lambda + \frac{2(v_1 - v_m)}{h^2} \right) x.$$

Thus

$$\lambda + \frac{2(v_1 - v_m)}{h^2} \leq \lambda$$

is an eigenvalue of $D^2(u^h; v_m)$, with equality only if $v_1 = v_m$. In particular, the eigenvalues of the discrete Hessian are decreasing functions of the root v .

Since by assumption, the discrete Monge-Ampère equation has a root that yields a positive-definite Hessian, at least one of the roots v_1, v_2, v_3 will yield three positive eigenvalues. As the eigenvalues are decreasing in v , the smallest root v_1 must yield three positive eigenvalues.

Suppose now that another root $v_m > v_1$ also produces positive eigenvalues. Since both v_1, v_m satisfy the discrete Monge-Ampère equation, we have

$$f = \lambda_1(D^2(u^h; v_m))\lambda_2(D^2(u^h; v_m))\lambda_3(D^2(u^h; v_m)) < \lambda_1(D^2(u^h; v_1))\lambda_2(D^2(u^h; v_1))\lambda_3(D^2(u^h; v_1)) = f.$$

This is a contradiction, which means that the smallest real root v_1 is the only root that yields three positive eigenvalues.

We conclude that the smallest real root of the cubic is the root that corresponds to the convex solution of the discrete Monge-Ampère equation (2.7). \square

3 FAS multigrid method

The drawback of directly using the nonlinear Gauss-Seidel iteration is that the number of iterations required for convergence increases with the number of discretization points, so the total solution time grows super-linearly with the number of variables. However, the nonlinear Gauss-Seidel iteration can be used as an effective smoother, which makes it particularly effective when combined with the nonlinear multigrid method that follows.

In this section, we introduce the standard nonlinear FAS multigrid method for solving our discretized nonlinear problem. As a highly efficient iterative algorithm, the FAS multigrid method is a nonlinear generalization of the linear multigrid algorithm, which typically has optimal computational complexity for linear elliptic PDEs. The FAS multigrid method provides a powerful approach for handling nonlinear equations without the need for the global linearization required by Newton's method. Unlike with Newton's method, there is typically no need to initialize the solver with a very good initial guess. Below, we briefly describe the FAS multigrid algorithm as well as its FMG version based on the standard textbooks [BL11, BHM00, TOS01, Saa03].

For a general nonlinear system that is discretized using the fine mesh-size h

$$S_h(w_h) = b_h,$$

one V-cycle FAS multigrid iteration is recursively defined in Algorithm 1 [BHM00, TOS01, Saa03].

Algorithm 1	FAS multigrid V-cycle iteration (with $H = 2h$)
Steps	$w_h := \text{FAS}(h, S_h, w_h^0, b_h)$
–	IF ($h == h_0$)
(1)	Approximately solve: $S_{h_0}(w_{h_0}) = b_{h_0}$
–	ELSE
(2)	Pre-smooth ν_1 times: $w_h := \text{SMOOTH}^{\nu_1}(S_h, w_h^0, b_h)$
(3)	Restriction residual: $r_H := I_h^H(b_h - S_h(w_h))$
(4)	Initialize coarse guess: $u_H := I_h^H w_h, w_H := \tilde{I}_h^H w_h$
(5)	Define coarse r.h.s.: $b_H := S_H(w_H) + r_H$
(6)	Recursion: $u_H := \text{FAS}(H, S_H, u_H, b_H)$
(7)	Prolongation: $\delta_h := I_H^h(u_H - w_H)$
(8)	Correction: $w_h := w_h + \delta_h$
(9)	Post-smooth ν_2 times: $w_h := \text{SMOOTH}^{\nu_2}(S_h, w_h, b_h)$
–	ENDIF
(10)	RETURN w_h .

In a single FAS V-cycle iteration, the fine-grid solution first undergoes a small number of smoothing iterations. Following this, the solution and residual are restricted to a coarser grid ($H = 2h$) using two (possibly different) restriction operators (I_h^H, \tilde{I}_h^H). A new V-cycle iteration is then performed on this coarser level, with this procedure proceeding recursively until the grid reaches the coarsest level $h_0 \gg h$.

At the coarsest level, the underlying problem size has become so small that it can be (approximately) solved using a very small number of Gauss-Seidel iterations.

Computing the solution on the coarser level H leads to a coarse approximation of the solution. A prolongation operator (I_H^h) transfers this approximation to the fine grid, which provides a coarse grid correction in the fine grid solution. The fine grid solution is further corrected using a small number of smoothing iterations.

As suggested in [TOS01], the approximation restriction operator \tilde{I}_h^H is chosen as straight injection. Depending on the application, there are many possible choices for I_h^H and I_H^h . For a better computational efficiency, we choose to use the residual restriction operator I_h^H from half-weighting averaging with the following stencil form

$$I_h^H = \frac{1}{12} \left[\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}^H \begin{bmatrix} 0 & 1 & 0 \\ 1 & 6 & 1 \\ 0 & 1 & 0 \end{bmatrix}^H \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}^H \right]_h$$

and the prolongation operator I_H^h from trilinear interpolation with a corresponding stencil form

$$I_H^h = \frac{1}{64} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}^h \begin{bmatrix} 2 & 4 & 2 \\ 4 & 8 & 4 \\ 2 & 4 & 2 \end{bmatrix}^h \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}^h \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}^H$$

Other types of I_h^H and I_H^h are also available when better accuracy is required.

We also note that a straightforward definition of the coarse grid operator S_H is possible by simply using the discretized equations with a coarser step size H . This is the most common choice for a fully structured mesh. We point out that the so-called Galerkin coarse grid operator, which constitutes a core component of the popular linear algebraic multigrid, is not easily used with the nonlinear FAS multigrid framework [TOS01].

The last but most crucial component of the method is an effective smoother SMOOTH, whose major function is to eliminate the high-frequency components of the approximation errors. As a standalone solver, the chosen smoothing iteration may converge very slowly as the mesh refines. This is the case for the nonlinear Gauss-Seidel iteration developed in the previous section. However, because it damps the high frequency components of the error, it will serve as an effective smoother SMOOTH in Algorithm 1. Such a smoothing property is numerically illustrated in Figure 1, which shows that two iterations seem to be sufficient to dramatically smooth out high-frequency errors. This is unsurprising given that the classical linear Gauss-Seidel iteration has been widely recognized as a benchmark smoother in the linear multigrid method. In practice, one or two pre- and post-smoothing iterations typically give the best overall performance. More pre- and post-smoothing iterations could be helpful when the smoothing effect of the smoother is weak or degraded by the possible singular solutions.

The efficiency of the above FAS multigrid iterative solver can be further improved when its initial guess is derived from the idea of nested iteration, which leads to the most efficient full multigrid (FMG)




Figure 1: The smoothing effect of our nonlinear Gauss-Seidel iterations with Example 1. Left: the initial oscillating error is chosen to be $\sin(\pi x) \sin(\pi y) \sin(\pi z) + 0.5 \sin(15\pi x) \sin(15\pi y) \sin(15\pi z)$ with a maximum norm 1.4712. Right: after 2 Gauss-Seidel iterations, the approximation error becomes smooth with a maximum norm 0.7538. Both error surface plots are only snapshots at $z = 1/2$ for better visualization.

algorithm [BHM00, TOS01, Saa03]. The FMG algorithm based on above the FAS V-cycle multigrid iteration is recursively defined in Algorithm 2.

Algorithm 2		FMG-FAS algorithm (with $H = 2h$)
Steps		$w_h := \text{FMG}(h, S_h, b_h)$
–	IF ($h == h_0$)	
(1)	Approximately solve:	$S_{h_0}(w_{h_0}) = b_{h_0}$
–	ELSE	
(2)	Restriction:	$b_H := \tilde{I}_h^H b_h$
(3)	Recursion:	$w_H := \text{FMG}(H, S_H, b_H)$
(4)	Cubic interpolation:	$w_h := \hat{I}_H^h w_H$
(5)	FAS V-cycle iteration:	$w_h := \text{FAS}(h, S_h, w_h, b_h)$
–	ENDIF	
(6)	RETURN w_h .	

In Algorithm 2, no initial guess is required at the finest level h , but we do need an initial guess at the coarsest level h_0 if we are going to solve the coarsest nonlinear system $S_{h_0}(w_{h_0}) = b_{h_0}$ using our Gauss-Seidel iteration. Fortunately, the coarsest nonlinear system has a very small dimension with effectively one unknown on a $3 \times 3 \times 3$ grid, the remaining values being determined by the Dirichlet boundary conditions. Thus it is easy to obtain a good approximate solution no matter what initial guess is used. In our implementation, we simply take the initial guess at the coarsest level h_0 to be identically zero. Often, a cubic interpolation \hat{I}_H^h will be used in the above FMG algorithm so that one full FMG iteration can deliver an approximation with the desired discretization accuracy. This is useful if the underlying FAS V-cycle multigrid iteration converges satisfactorily, as is the case for Poisson equation solvers [TOS01]. Due to the strong nonlinearity of our problem, we do not expect one FMG sweep to be sufficient. However, it

does provide a very good initial guess for the FAS V-cycle multigrid iterations. In both algorithms, the coarse level problem is solved approximately by performing just one nonlinear Gauss-Seidel iteration.

4 Numerical examples

In this section, we test our FAS multigrid algorithm using several examples available in the literature. All simulations are implemented using MATLAB on a laptop PC with Intel(R) Core(TM) i3-3120M CPU@2.50GHz and 12GB RAM. We apply one FMG iteration as an initialization step. Due to the nonlinearity, this initialization step is necessary for ensuring mesh-independent convergence of the FAS algorithm. For each FAS V-cycle, we perform only two pre-smoothing and two post-smoothing Gauss-Seidel iterations. The coarsest level ($h_0 = 1/2$) problem is approximately solved using only one Gauss-Seidel iteration. Numerical simulations indicate that the FMG-FAS algorithm has almost the same efficiency using more accurate solvers at the coarsest level.

At the l -th iteration, let u_h^l denote the computed approximation of solution and r_h^l be the corresponding residual vector. We use the pre-specified reduction in the relative residuals as the stopping criterion

$$\frac{\|r_h^l\|_2}{\|r_h^0\|_2} \leq 10^{-6},$$

where $\|\cdot\|_2$ denotes the discrete L_2 norm. Our numerical simulations show that the chosen tolerance 10^{-6} is sufficient to achieve the level of discretization error. Let u_h denote the computed solution approximated at the last iteration. We first compute the infinity norm of the approximation error

$$\text{Error}(h) = \|u_h - u\|_\infty$$

and then estimate the experimental order of accuracy by computing the logarithmic ratios of the approximation errors between two successive refined meshes, i.e.,

$$\text{Order} = \log_2 \left(\frac{\text{Error}(2h)}{\text{Error}(h)} \right),$$

which should be close to two provided that the scheme delivers a second-order accuracy.

In most cases, only one FMG iteration is needed for convergence. In special cases, we may need a few extra multigrid V-cycle iterations to fulfill the above convergence criterion. In 3D problems, using two pre- and post-smoothing steps, the computational cost of one FAS V-cycle is about 32/7 the cost of a single Gauss-Seidel iteration on the finest mesh. One FMG-FAS iteration costs roughly 5 times the cost of a single Gauss-Seidel iteration on the finest mesh if we ignore the additional cost of restriction and interpolation.

When we use the nonlinear Gauss-Seidel iterations as a standalone solver, it is challenging to pick a good initial guess such that the iteration converges quickly. For the purpose of simple comparison, we always choose the slightly perturbed solution $1.01u$ as the initial guess for the nonlinear Gauss-Seidel solver, where u is the analytic solution. Even using this impractical initial guess, the Gauss-Seidel solver

cannot compete with the efficiency of the FMG-FAS method.

4.1 Two dimensional examples

We begin by applying the corresponding 2D version of our FMG-FAS method using the nonlinear Gauss-Seidel iteration developed in [BFO10] as the smoother. The scheme involves a 9-point stencil, and we test the Gauss-Seidel smoother using the red-black ordering and the plain lexicographic ordering, respectively. The FMG-FAS multigrid method delivers a far better computational efficiency than the results of [BFO10], which used the Gauss-Seidel iteration as a standalone solver. In particular, the computation time scales almost linearly with respect to the degrees of freedom, which is verified by a consistent four-fold increase in the CPU time as the stepsize h is halved.

Example 1. Let $\Omega = (0, 1)^2$. Choose

$$f(x, y) = (1 + x^2 + y^2) \exp(x^2 + y^2)$$

such that an analytic solution is

$$u(x, y) = \exp\left(\frac{x^2 + y^2}{2}\right).$$

Example 2. Let $\Omega = (0, 1)^2$. Choose

$$f(x, y) = \frac{1}{\sqrt{x^2 + y^2}}$$

such that an exact solution is

$$u(x, y) = \frac{2\sqrt{2}}{3}(x^2 + y^2)^{3/4}.$$

Example 3. Let $\Omega = (0, 1)^2$. Choose

$$f(x, y) = \frac{2}{(2 - x^2 - y^2)^2}$$

such that an exact solution is

$$u(x, y) = -\sqrt{2 - x^2 - y^2}.$$

Table 1: Results for Example 1, 2, and 3 with the FMG-FAS algorithm (red-black ordering).

N	Example 1				Example 2				Example 3			
	Error	Order	Iter	CPU	Error	Order	Iter	CPU	Error	Order	Iter	CPU
128	4.5e-06	–	1	0.1	3.8e-05	–	1	0.1	8.9e-03	–	2	0.1
256	1.1e-06	2.0	1	0.1	1.3e-05	1.5	1	0.1	6.4e-03	0.5	2	0.2
512	2.8e-07	2.0	1	0.3	4.7e-06	1.5	1	0.3	4.5e-03	0.5	1	0.3
1024	7.0e-08	2.0	1	1.0	1.7e-06	1.5	1	1.1	3.2e-03	0.5	1	1.1
2048	1.7e-08	2.0	1	3.6	5.9e-07	1.5	1	4.0	2.3e-03	0.5	1	3.8
4096	4.3e-09	2.0	1	16.2	2.1e-07	1.5	1	16.5	1.6e-03	0.5	1	18.0
8192	1.0e-09	2.1	1	67.1	7.4e-08	1.5	1	69.1	1.1e-03	0.5	1	70.3

Table 2: Results for Example 1, 2, and 3 with the FMG-FAS algorithm (lexicographic ordering).

	Example 1				Example 2				Example 3			
N	Error	Order	Iter	CPU	Error	Order	Iter	CPU	Error	Order	Iter	CPU
128	5.1e-06	–	1	0.1	4.3e-05	–	1	0.1	8.9e-03	–	3	0.1
256	1.3e-06	2.0	1	0.2	1.5e-05	1.5	1	0.3	6.3e-03	0.5	2	0.2
512	3.3e-07	2.0	1	0.8	5.4e-06	1.5	1	0.9	4.4e-03	0.5	1	0.8
1024	8.2e-08	2.0	1	3.2	1.9e-06	1.5	1	3.27	3.2e-03	0.5	1	3.1
2048	2.1e-08	2.0	1	12.7	6.8e-07	1.5	1	12.9	2.2e-03	0.5	1	12.6
4096	5.1e-09	2.0	1	53.2	2.4e-07	1.5	1	54.6	1.6e-03	0.5	1	53.5
8192	1.3e-09	2.0	1	245.6	8.5e-08	1.5	1	259.4	1.1e-03	0.5	1	250.7

For simplicity, we combine the results of Examples 1, 2, and 3 in Tables 1 and 2. The CPU times manifest the desired linear time complexity of our FMG-FAS algorithm. The non-smoothness of the solution in Examples 2 and 3 degrade the second-order accuracy of the approximations to $O(h^{3/2})$ and $O(h^{1/2})$, respectively. Our algorithm seems to be able to effectively and efficiently handle this type of mildly singular solution. It also shows that the red-black ordering has slightly better approximation accuracy as well as superior computational efficiency due to the faster vectorization in MATLAB implementation. According to our numerical experiments, for strongly singular solutions (not even in $H^2(\Omega)$), additional FAS V-cycle iterations may be needed as the mesh is refined. In that case, the deterioration of our algorithm is associated with the possible failure of the corresponding finite difference approximations, which are not provable convergent, and can break down in the singular case.

4.2 Three dimensional examples

Next we turn our attention to the three-dimensional Monge-Ampère equation using the Gauss-Seidel iteration introduced in this article.

Example 4 [FO11a]. Let $\Omega = (0, 1)^3$. Choose

$$f = (1 + x^2 + y^2 + z^2) \exp\left(\frac{3(x^2 + y^2 + z^2)}{2}\right)$$

such that the exact solution is

$$u = \exp\left(\frac{x^2 + y^2 + z^2}{2}\right).$$

In Table 3 we report the relative residual norms, infinity norm of errors, iteration numbers, and CPU times of the FMG-FAS algorithm and the Gauss-Seidel solver, respectively. The ‘Error’ column and ‘Order’ column show that the central finite difference discretization indeed achieves the expected second-order accuracy in the maximum norm. The ‘Iter’ column implies that our FMG-FAS algorithm possesses a mesh-independent convergence. More specifically, here ‘Iter’=1 implies that only one FMG iteration is sufficient to fulfill our stopping criteria; no further V-cycles are necessary. The achieved relative residual norms (column ‘RelRes’) turns out to be much lower, thanks to the fast convergence of the FMG algorithm. We can do a simple calculation to see that one FMG iteration indeed costs about 5 Gauss-Seidel iterations. Consider the mesh size $N = 32$, one Gauss-Seidel iteration takes about $324.6/867 \approx 0.4$ sec-

ond and hence one FMG iteration should require approximately $5 \times 0.4 = 2.0$ seconds, which is almost exactly what we achieved in our algorithm. The ‘CPU’ column indicates that our FMG-FAS algorithm has a roughly linear time complexity since the CPU time increases by eight times as the mesh size $h = 1/N$ is halved.

On the other hand, the nonlinear Gauss-Seidel iterations as a standalone solver requires an increasing number of iterations as the mesh size decreases, which is expected. Even starting with an artificial initial guess that is impractically close the desired true solution, the nonlinear Gauss-Seidel iteration still does not produce sufficiently accurate approximations using the same stopping criterion (compare the row with $N = 32$). It is also interesting to notice that the relative residuals (column ‘RelRes’) of the Gauss-Seidel solver are much larger than that of the FMG-FAS algorithm.

For comparison, Table 4 also lists the results reported in [FO11b] and [BN12], where a wide-stencil hybrid finite difference solver and a finite element approximation were used for the same test problem. Though we can not fairly compare the CPU times directly as they were computed on different machines, it is clear that the FMG-FAS method is the only one the scales linearly with the number of degrees of freedom. The accuracy of the FMG-FAS approach is also competitive. As a whole, our FMG-FAS method appears superior than both the wide-stencil finite difference solvers and the finite element approximations for this type of simple problem with a sufficiently smooth solution.

We also mention that when $N = 128$, the second order accuracy is no longer observed. This is due to the limitations of machine precision and round-off errors during the computation. We note that the discretized Monge-Ampère operator involves division by $h^6 \approx 2 \times 10^{-13}$ when $N = 128$. We could get around this using higher precision arithmetic. However, the global error is now 6×10^{-6} , which is small. The key is that we are able to resolve the data, which necessitates solving on large grids, and the resulting accuracy is sufficient for the applications we have in mind.

Table 3: Results for Example 4 with the FMG-FAS algorithm and the Gauss-Seidel solver.

N	FMG-FAS					Gauss-Seidel				
	RelRes	Error	Order	Iter	CPU	RelRes	Error	Order	Iter	CPU
8	3.1e-07	1.2e-03	–	1	0.1	9.3e-07	1.2e-03	–	72	0.4
16	9.5e-09	2.7e-04	2.1	1	0.2	9.9e-07	3.0e-04	2.0	270	11.5
32	2.1e-10	5.7e-05	2.3	1	2.0	1.0e-06	8.1e-05	1.9	867	324.6
64	4.1e-12	1.4e-05	2.0	1	16.9					
128	5.2e-13	6.2e-06	1.2	1	136.9					

Table 4: Reported results for Example 4 from [FO11b] and [BN12].

Wide-stencil finite difference [FO11b]				Finite element [BN12]			
N	DOFs	Error	CPU	h	DOFs	Error	CPU
8	343	1.51e-02	0.04	1/4	1581	9.12e-03	3.89
12	1331	1.40e-02	0.10	1/8	12611	2.29e-03	38.55
16	3375	1.29e-02	0.71	1/12	42798	4.49e-04	140.89
22	9261	1.21e-02	6.72	1/16	99436	2.73e-04	355.78
32	29791	1.11e-02	86.63	1/20	195110	1.21e-04	803.47

Example 5. Let $\Omega = (0, \pi)^3$. Choose

$$f = (\sin(x) + 1)(\sin(y) + 1)(\sin(z) + 1)$$

such that an exact solution is

$$u = -\sin(x) - \sin(y) - \sin(z) + (x^2 + y^2 + z^2)/2.$$

In Table 5 we report the relative residual norms, infinity norm of errors, iteration numbers, and CPU times of the FMG-FAS algorithm and the Gauss-Seidel solver, respectively. Again, we observed the similar excellent performance of our FMG-FAS algorithm against with the nonlinear Gauss-Seidel solver.

Table 5: Results for Example 5 with the FMG-FAS algorithm and the Gauss-Seidel solver.

N	FMG-FAS					Gauss-Seidel				
	RelRes	Error	Order	Iter	CPU	RelRes	Error	Order	Iter	CPU
8	6.8e-08	1.7e-02	–	2	0.1	9.2e-07	1.7e-02	–	82	0.4
16	2.1e-08	3.3e-03	2.4	1	0.2	9.8e-07	4.3e-03	2.0	296	12.9
32	2.7e-10	8.1e-04	2.0	1	2.0	1.0e-06	1.1e-03	2.0	954	356.2
64	3.7e-12	2.1e-04	2.0	1	16.4					
128	4.8e-13	5.2e-05	2.0	1	134.1					

Example 6 [BN12]. Let $\Omega = (0, 1)^3$. Choose

$$f = (1/16)(x^2 + y^2 + z^2)^{-3/4}$$

such that an analytic solution is

$$u = (1/3)(x^2 + y^2 + z^2)^{3/4}.$$

In Table 6 we report the relative residual norms, infinity norm of errors, iteration numbers, and CPU times of the FMG-FAS algorithm and the Gauss-Seidel solver, respectively. In contrast to the previous two examples, the solution to this problem is not smooth as there is a singularity at the origin. Nevertheless, our central finite difference scheme still approximates the solutions with a reasonable $O(h^{3/2})$ accuracy in the infinity norm. Similarly, our FMG-FAS algorithm greatly outperforms the nonlinear Gauss-Seidel solver.

Table 6: Results for Example 6 with the FMG-FAS algorithm and the Gauss-Seidel solver.

N	FMG-FAS					Gauss-Seidel				
	RelRes	Error	Order	Iter	CPU	RelRes	Error	Order	Iter	CPU
8	2.7e-07	5.3e-04	–	1	0.1	9.7e-07	4.0e-04	–	86	0.4
16	1.1e-08	2.1e-04	1.4	1	0.2	9.9e-07	1.4e-04	1.5	312	13.2
32	4.8e-10	7.4e-05	1.5	1	2.0	9.9e-07	5.1e-05	1.5	925	345.8
64	2.0e-11	2.6e-05	1.5	1	16.4					
128	1.2e-12	9.3e-06	1.5	1	133.9					

Example 7 [FO11a]. Let $\Omega = (0, 1)^3$. Choose

$$f(x, y) = 3(3 - x^2 - y^2 - z^2)^{-5/2}$$

such that an analytic solution is

$$u = -\sqrt{3 - x^2 - y^2 - z^2}.$$

In Table 7 we report the relative residual norms, infinity norm of errors, iteration numbers, and CPU times of the FMG-FAS algorithm and the Gauss-Seidel solver, respectively. This example is more challenging due to the unbounded gradient at the boundary. The ‘Error’ column shows our scheme has a roughly $O(h^{1/2})$ accuracy in the infinity norm, which is consistent with the error in standard discretizations for the two-dimensional version of this example [BFO10]. Obviously, our FMG-FAS algorithm requires significantly less CPU time than the Gauss-Seidel solver. We note that the slightly better accuracy achieved by the Gauss-Seidel solver is due to the unrealistic initial guess used.

Table 7: Results for Example 7 with the FMG-FAS algorithm and the Gauss-Seidel solver.

N	FMG-FAS					Gauss-Seidel				
	RelRes	Error	Order	Iter	CPU	RelRes	Error	Order	Iter	CPU
8	1.7e-07	1.4e-03	–	1	0.1	8.9e-07	1.4e-03	–	75	0.4
16	2.1e-08	1.4e-03	0.0	1	0.2	9.7e-07	1.4e-03	0.0	266	12.2
32	2.7e-08	4.3e-03	-1.7	1	2.0	9.9e-07	1.1e-03	0.3	854	320.2
64	4.9e-09	2.2e-03	1.0	1	16.3					
128	1.7e-09	1.6e-03	0.4	1	133.1					

5 Conclusions

In this paper, we presented a fast method for solving the three-dimensional elliptic Monge-Ampère equation, focusing on the Dirichlet problem and smooth solutions. We combined a nonlinear Gauss-Seidel iterative method with a standard centered difference discretization. The Gauss-Seidel iteration was then used as an effective smoother, in combination with a nonlinear full approximation scheme (FAS) multigrid method. Two and three dimensional numerical examples were computed to demonstrate the computational efficiency of the proposed multigrid method, and comparison was made with other available approaches. In particular, we show that the computational cost of the method scales approximately linearly. Computations on a recent laptop allowed for grid sizes of 128^3 in less than three minutes. More importantly, since the methods scale well, implementation on more powerful computers will lead to a feasible approach to solving problems with the larger resolutions that occur in, for example three dimensional image registration

Future work will be to extend these ideas to filtered almost-monotone finite differences schemes in order to build solvers that apply to the singular solutions that can occur in applications. We would also like to extend the method to the Optimal Transportation problem with applications to three dimensional image registration, or to parameter identification.

References

- [BFO10] Jean-David Benamou, Brittany D. Froese, and Adam M. Oberman. Two numerical methods for the elliptic Monge-Ampère equation. *M2AN Math. Model. Numer. Anal.*, 44(4):737–758, 2010.
- [BFO14] Jean-David Benamou, Brittany D. Froese, and Adam M. Oberman. Numerical solution of the optimal transportation problem using the Monge-Ampère equation. *J. Comput. Phys.*, 260:107–126, 2014.
- [BGNS11] Susanne C. Brenner, Thirupathi Gudi, Michael Neilan, and Li-yeng Sung. C^0 penalty methods for the fully nonlinear Monge-Ampère equation. *Math. Comp.*, 80(276):1979–1995, 2011.
- [BHM00] William L. Briggs, Van Emden Henson, and Steve F. McCormick. *A multigrid tutorial*. SIAM, Philadelphia, PA, 2000.
- [BL11] A. Brandt and O.E. Livne. *Multigrid Techniques: 1984 Guide with Applications to Fluid Dynamics, Revised Edition*. Classics in Applied Mathematics. SIAM, 2011.
- [BN12] Susanne Cecelia Brenner and Michael Neilan. Finite element approximations of the three dimensional Monge-Ampère equation. *ESAIM Math. Model. Numer. Anal.*, 46(5):979–1001, 2012.
- [Bra77] Achi Brandt. Multi-level adaptive solutions to boundary-value problems. *Math. Comp.*, 31(138):333–390, 1977.
- [BS91] Guy Barles and Panagiotis E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Anal.*, 4(3):271–283, 1991.
- [CC11] Noppadol Chumchob and Ke Chen. A robust multigrid approach for variational image registration models. *J. Comput. Appl. Math.*, 236(5):653–674, 2011.
- [CDF11] L. Chacón, G. L. Delzanno, and J. M. Finn. Robust, multidimensional mesh-motion based on Monge-Kantorovich equidistribution. *J. Comput. Phys.*, 230(1):87–103, 2011.
- [CNS84] L. Caffarelli, L. Nirenberg, and J. Spruck. The Dirichlet problem for nonlinear second-order elliptic equations. I. Monge-Ampère equation. *Comm. Pure Appl. Math.*, 37(3):369–402, 1984.
- [DCF⁺08] Gian Luca Delzanno, Luis Chacón, John M Finn, Y Chung, and Giovanni Lapenta. An optimal robust equidistribution method for two-dimensional grid adaptation based on Monge-Kantorovich optimization. *Journal of Computational Physics*, 227(23):9841–9864, 2008.

- [EF14] Björn Engquist and Brittany D Froese. Application of the Wasserstein metric to seismic signals. *Communications in Mathematical Sciences*, 12(5), 2014.
- [FGN13] Xiaobing Feng, Roland Glowinski, and Michael Neilan. Recent developments in numerical methods for fully nonlinear second order partial differential equations. *SIAM Rev.*, 55(2):205–267, 2013.
- [FN09] Xiaobing Feng and Michael Neilan. Vanishing moment method and moment solutions for fully nonlinear second order partial differential equations. *Journal of Scientific Computing*, 38(1):74–98, 2009.
- [FO11a] Brittany D. Froese and Adam M. Oberman. Convergent finite difference solvers for viscosity solutions of the elliptic Monge-Ampère equation in dimensions two and higher. *SIAM J. Numer. Anal.*, 49(4):1692–1714, 2011.
- [FO11b] Brittany D. Froese and Adam M. Oberman. Fast finite difference solvers for singular solutions of the elliptic Monge-Ampère equation. *J. Comput. Phys.*, 230(3):818–834, 2011.
- [FO12] Brittany D. Froese and Adam M. Oberman. Convergent almost monotone finite difference approximations for viscosity solutions of the elliptic Monge-Ampère partial differential equation. *submitted*, 2012.
- [GZ00] Murli M. Gupta and Jun Zhang. High accuracy multigrid solution of the 3D convection-diffusion equation. *Appl. Math. Comput.*, 113(2-3):249–274, 2000.
- [HM07] Eldad Haber and Jan Modersitzki. Image registration with guaranteed displacement regularity. *International Journal of Computer Vision*, 71(3):361–372, 2007.
- [HPM⁺09] Eldad Haber, Gallagher Pryor, John Melonakos, Allen Tannenbaum, et al. 3d nonrigid registration via optimal mass transport on the GPU. *Medical image analysis*, 13(6):931–940, 2009.
- [HRT10] Eldad Haber, Tauseef Rehman, and Allen Tannenbaum. An efficient numerical method for the solution of the L_2 optimal mass transfer problem. *SIAM Journal on Scientific Computing*, 32(1):197–211, 2010.
- [HZTA04] Steven Haker, Lei Zhu, Allen Tannenbaum, and Sigurd Angenent. Optimal mass transport for registration and warping. *International Journal of Computer Vision*, 60(3):225–240, 2004.
- [LeV07] Randall LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*. SIAM, Philadelphia, PA, USA, 2007.
- [LR05] Grégoire Loeper and Francesca Rapetti. Numerical solution of the Monge-Ampère equation by a Newton’s algorithm. *C. R. Math. Acad. Sci. Paris*, 340(4):319–324, 2005.

- [Nei13] Michael Neilan. Quadratic finite element approximations of the Monge-Ampère equation. *Journal of Scientific Computing*, 54(1):200–226, 2013.
- [Obe08] Adam M. Oberman. Wide stencil finite difference schemes for the elliptic Monge-Ampère equation and functions of the eigenvalues of the Hessian. *Discrete Contin. Dyn. Syst. Ser. B*, 10(1):221–238, 2008.
- [OR00] J. M. Ortega and W. C. Rheinboldt. *Iterative solution of nonlinear equations in several variables*. SIAM, Philadelphia, PA, 2000.
- [PTTBvR⁺14] C. R. Prins, J. H. M. Ten Thije Boonkamp, J. van Roosmalen, W. L. Ijzerman, and T. W. Tukker. A Monge-Ampère-solver for free-form reflector design. *SIAM J. Sci. Comput.*, 36(3):B640–B660, 2014.
- [Saa03] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, Philadelphia, PA, 2003.
- [SAK13] Louis-Philippe Saumier, Martial Agueh, and Boualem Khouider. An efficient numerical algorithm for the L2 optimal transport problem with periodic densities. *IMA Journal of Applied Mathematics*, page hxt032, 2013.
- [Soi11] Ashmeen Kaur Soin. Multigrid for Elliptic Monge-Ampère Equation. Master’s thesis, University of Waterloo, Waterloo, Ontario, Canada, 2011.
- [SWR11] Mohamed M. Sulman, J. F. Williams, and Robert D. Russell. An efficient approach for the numerical solution of the Monge-Ampère equation. *Appl. Numer. Math.*, 61(3):298–307, 2011.
- [Tad12] Eitan Tadmor. A review of numerical methods for nonlinear partial differential equations. *Bull. Amer. Math. Soc. (N.S.)*, 49(4):507–554, 2012.
- [TOS01] U. Trottenberg, C. W. Oosterlee, and A. Schüller. *Multigrid*. Academic Press Inc., San Diego, CA, 2001.
- [TW08] Neil S Trudinger and Xu-Jia Wang. The Monge-Ampere equation and its geometric applications. *Handbook of geometric analysis*, 1:467–524, 2008.
- [Zha98] Jun Zhang. Fast and high accuracy multigrid solution of the three-dimensional Poisson equation. *J. Comput. Phys.*, 143(2):449–461, 1998.
- [ZPF10] V. Zheligovsky, O. Podvigina, and U. Frisch. The Monge-Ampère equation: Various forms and numerical solution. *Journal of Computational Physics*, 229(13):5043–5061, 2010.